

Jisc is the UK's expert body for digital technology and digital resources in higher education, further education and research. Since its foundation in the early 1990s, Jisc has played a pivotal role in the adoption of information technology by UK universities and colleges, supporting them to improve learning, teaching, the student experience and institutional efficiency, as well as enabling more powerful research.

Jisc operates the UK's high-speed National Research and Education Network, Janet, connecting UK universities, colleges, research organisations and schools to each other and to the Internet. To facilitate research and education, Janet is designed to be as open and neutral as possible. For many years, the network and its customers have been the target of large denial of service attacks. Our incident response team, Janet CSIRT, uses temporary filters to mitigate the impact of these and other attacks. However we also follow Internet best practice by permanently filtering outbound spoofed traffic to protect others from the consequences of attacks on systems connected to our network.

Summary

We welcome BEREC's recognition of the importance of filtering to protect the security of internet services and users. However, while most of types of filtering identified in [paragraph 80 of the draft guidelines](#) can be implemented in response **to** a particular threat to a network, this is not true of filtering to protect other networks **from** threats created by the filtering network's own users.

BEREC's draft guidelines identify one such class of filtering – spoofed addresses. This type of filtering can, as discussed below, only be done by the networks that originate traffic. The Internet Engineering Task Force has long considered it Best Current Practice against denial of service attacks, as documented in their [BCP-38](#). This recommends that all networks be permanently configured to detect and block packets with spoofed source addresses, before they leave the originating network. This recommendation is promoted by network operators (for example the FENIX group in the Czech Republic) and regulators (for example FICORA in Finland). We are concerned that the draft regulations, by stating that permanent filtering should be considered a breach of network neutrality, would seriously harm these efforts to protect the security and stability of networks and services.

Permanent filtering of spoofed addresses is not only an effective way to reduce the opportunity to conduct denial of service attacks, it also distinguishes very precisely between legitimate and non-legitimate traffic. Unlike other types of security filtering it should not, therefore, affect network neutrality in any way. The only packets that will be blocked are those that, either accidentally or deliberately, do not conform to the fundamental Internet Protocol standard. Computers sending these packets would not, even on an unfiltered network, receive any internet service, since the response packets would never reach them. Filtering spoofed packets will have no effect on the computers sending those packets and only beneficial effects on the rest of the network.

We therefore encourage BEREC to recognise this type of filtering as not constituting a breach of network neutrality.

Discussion

Most Internet denial of service attacks use a technique known as amplification (see Arbor Networks, "Worldwide Infrastructure Security Report, Volume XI (2015)", page 24). This has been compared to the attacker asking for a mail order catalogue to be sent to the victim: by sending small postcards to a legitimate third party the attacker can create a much greater load on the victim's mail delivery service.

The Internet version of the technique likewise involves an attacker sending a small message to a third party that causes that third party to send a much larger message to the victim. The Network Time Protocol (NTP) can generate responses 500 times larger than the request, [many other services provide amplification factors of more than 100](#). Most attackers use compromised computers to send their request packets, obtaining a further level of amplification. A single command sent to a few hundred compromised computers, each of which generates amplification requests at the speed of a typical ADSL connection, can generate flows of tens or hundreds of gigabytes per second to the chosen victim. This is sufficient to fill the connection of almost any organisation, resulting in the victim's website and other services becoming inaccessible to legitimate visitors. Such attacks may be used, for example, for blackmail, activism, online gaming advantage, or to distract the victim from other hostile activity. By congesting other networks, their side-effects can cause instabilities across a wide area.

Amplification attacks are particularly hard for the victim to deal with, as the packets they receive are completely normal and come from legitimate sources. Any filtering that the victim or their network provider can implement in response to the attack will inevitably block legitimate traffic as well as the packets forming the attack. Similarly the services that are used for the amplification receive apparently normal requests, though perhaps at an increased rate, and respond in the normal way.

The only parties that can distinguish the packets involved in an attack are the networks that connect the compromised computers, controlled by the attacker, to the Internet. For the responses to be sent to the victim, the request packets must appear to come from the victim. Request packets are therefore sent with a "spoofed" source address – that of the intended victim – rather than the true addresses of the computers that generate them. This is the only point in the attack where abnormal packets are used, and where they can be accurately distinguished from legitimate traffic. Networks connecting users or organisations to the Internet should know which IP address ranges those users or organisations legitimately use: any packet that has a source address outside these ranges can be detected and blocked. Once packets reach transit providers the diverse connectivity of the Internet makes it practically impossible to distinguish those having a source address that does not match their network of origin.

Blocking spoofed packets was identified as Best Current Practice against denial of service attacks by the Internet Engineering Task Force in 2000. Wider adoption of [this recommendation, known as BCP-38](#), has been encouraged by many global and national campaigns, including the Internet Society's [Mutually Agreed Norms for Routing Security](#) (2014) and Czech Internet Service Providers' [FENIX project](#) (2013). The Finnish telecommunications regulator, FICORA, makes [BCP-38 implementation mandatory for ISPs in Finland](#).

Unlike filtering by the victim of a denial of service attack, address spoofing filters must be in place permanently. The volume of spoofed traffic from any individual computer is unlikely to

be sufficient to trigger its Internet Service Provider's alarms; other networks or services suffering from the attack cannot notify the source ISPs because the address spoofing prevents them identifying the source of the packets. BEREC's requirement that filters be enabled only in response to a particular threat is therefore likely to reduce (or at best slow down) the adoption of this important protection technique.

This would be a particularly unfortunate outcome of regulation designed to protect network neutrality, as filtering spoofed addresses is the most neutral technique available to prevent denial of service attacks. As discussed above, any filtering by the victim will inevitably also block legitimate packets, so will interfere with some genuine use of the network. By contrast, packets with spoofed source addresses can never form part of genuine use, because the responses will never reach the computer that originated them. Spoofed packets can only be created accidentally, through a misconfiguration of the sending computer, or maliciously. In the former case the computer will not receive services from the Internet whether or not its packets are filtered by its ISP. Such filtering therefore makes no difference to users' Internet experience and has no impact on network neutrality.